

Non-degeneracy of Pollard Rho Collisions

Stephen D. Miller* and Ramarathnam Venkatesan

August 31, 2008

Abstract

The Pollard ρ algorithm is a widely used algorithm for solving discrete logarithms on general cyclic groups, including elliptic curves. Recently the first nontrivial runtime estimates were provided for it, culminating in a sharp $O(\sqrt{n})$ bound for the collision time on a cyclic group of order n [4]. In this paper we show that for n satisfying a mild arithmetic condition, the collisions guaranteed by these results are nondegenerate with high probability: that is, the Pollard ρ algorithm successfully finds the discrete logarithm.

Keywords: Pollard Rho algorithm, discrete logarithm, random walk, expander graph, collision time, mixing time, spectral analysis.

1 Introduction

The Pollard ρ algorithm is, to date, the leading algorithm for solving discrete logarithm problems on general groups, including elliptic curves. The algorithm can be stated as follows. Let G be a cyclic group of order n generated by the element g ; n may assumed to be a large prime because of the Pohlig-Hellman reduction [1]. Let $h = g^y$ be the element whose discrete logarithm $y \neq 1$ (unknown) is to be found, and let $x_0 = h$ or a random power $g^{r_1}h^{r_2}$ (which turns out to be only slightly less general). Let $G = S_1 \cup S_2 \cup S_3$ be a random partition of G into three disjoint subsets, in which each element has

*Partially supported by NSF grant DMS-0601009 and an Alfred P. Sloan Foundation Fellowship.

a $1/3$ probability of belonging to each S_j .¹ Define an iteration $x_{k+1} = f(x_k)$, where

$$f(x) = \begin{cases} gx, & x \in S_1; \\ hx, & x \in S_2; \\ x^2, & x \in S_3. \end{cases} \quad (1.1)$$

At each stage x_k may be written as $g^{a_k y + b_k}$, where the coefficients $a_k, b_k \in \mathbb{Z}/n\mathbb{Z}$ are known. Iterate until a collision of values $x_k = x_\ell$ has been found, and if the collision is “non-degenerate” (meaning $(a_k, b_k) \neq (a_\ell, b_\ell)$), solve for the discrete logarithm using the formula $y = \frac{b_\ell - b_k}{a_k - a_\ell}$.

The algorithm is conjectured to run in time $O(\sqrt{n})$ with high probability. It is the only such algorithm which uses small memory and which works for general groups. Though faster algorithms are known for specific incarnations of cyclic groups², a theorem of Victor Shoup [7] asserts that no algorithm on a general group can be faster – aside from improving the implied multiplicative constant. For it to be successful, two things must happen:

1. A collision must be found in time $O(\sqrt{n})$.
2. This collision must be non-degenerate.

Item 1 has been the subject of a number of recent papers, before which there were no nontrivial bounds on the runtime at all. First, a collision time of $O(\sqrt{n}(\log n)^3)$ was shown in [5], which was successively improved by [3] and [4] to the optimal $O(\sqrt{n})$ bound.

The purpose of this paper is to address Item 2 for the Pollard ρ algorithm (it is, however, settled for some variants of Pollard ρ , as in [2]). Unfortunately as of yet we are unable to make the result unconditional, for it depends on the multiplicative order of 2 modulo n (the least positive integer k such that $2^k \equiv 1 \pmod{n}$). We prove the following result, which is a complete runtime analysis for almost all group orders n :

¹In practice, the assignment is accomplished using a hash function which is expected to behave randomly, as storing the partitions themselves would take up too much memory. A formal model would assume a cryptographically strong pseudo-random function whose underlying cryptographic primitive would have a security estimate exceeding the runtime of the Pollard Rho algorithm. Such implementation details to justify the random walk model used in our analysis are well understood.

²For example, index calculus provides a subexponential algorithm on the group \mathbb{F}_p^* , which is abstractly isomorphic to a cyclic group of order $n = p - 1$. Note that this is *not* itself an example of a prime order cyclic group as treated above: one must apply the Pohlig-Hellman reduction first.

1.2 Theorem. *Consider the Pollard ρ algorithm as above on a group $G = \langle g \rangle$ of prime order n , starting at a random point $x_0 = g^{r_1} h^{r_2}$. Suppose that the multiplicative order of 2 modulo n is at least $c_0(\log n)^3$, where c_0 is the absolute constant coming from Proposition 2.18. Then any Pollard ρ collision occurring before time T is nondegenerate with probability at least $1 - \frac{3T^2}{2n^2}$. In particular, the collisions guaranteed by [4] to occur with high probability within time $O(\sqrt{n})$ are nondegenerate with probability at least $1 - O(\frac{1}{n})$.*

Remarks: 1) Though the probability of nondegeneracy is heuristically higher than that of collisions, in practice it has been much more difficult to prove nondegeneracy.

2) The multiplicative order of 2 modulo n is typically quite large, e.g. it equals $n - 1$ if 2 generates $(\mathbb{Z}/n\mathbb{Z})^*$, which it frequently does. There do exist primes with multiplicative order the size of $\log n$ (e.g. Fermat and Mersenne primes), but those disobeying the condition in the theorem are quite rare. Indeed, we show in Lemma 3.3 that at most $O((\log X)^5)$ such primes p exist in the range $X \leq p \leq 2X$.

3) Even if 2 has small multiplicative order modulo n , there is always a prime $\ell = O((\log n)^6(\log \log n))$ which has multiplicative order at least $c_0(\log n)^3$. This is because a cyclic group has at most L^2 elements of order $\leq L$. (We thank the referee for supplying this argument.) If the Pollard ρ algorithm is modified to replace the squaring step by $x \mapsto x^\ell$ instead, the analysis here and in [3–5] applies and gives a completely rigorous proof of the same $O(\sqrt{n})$ runtime, with the same $1 - O(\frac{1}{n})$ success rate.

4) The reason we need to assume a random starting point, unlike in [5], is that we cannot rule out degeneracies in collisions occurring within the first few steps. Lemma 3.1, in particular, applies only to random starting points. Once the algorithm has proceeded for $c_0(\log n)^3$ steps a random point is reached regardless of the starting point, but we cannot guarantee a random position before then.

The strategy of the proof starts with the viewpoint that the Pollard ρ iteration can be modeled as a pseudo-random walk on the “Pollard ρ graph”: the graph whose vertices are elements of G , and whose (directed) edges have the form

$$x \longrightarrow xg, xh, \text{ or } x^2. \quad (1.3)$$

Indeed, until a collision occurs the iteration by (1.1) is in fact a random walk, because the destination from a vertex x depends only on its random assignment to one of the S_j ; however, it is important that the walk is no

longer random after this point, for it enters a loop. The coefficients $(a_k, b_k) \in (\mathbb{Z}/n\mathbb{Z})^2$ meanwhile likewise can be modeled as a random walk (until the time of collision) on the following “Pollard ρ coefficient graph”:

$$(a, b) \longrightarrow (a + 1, b), (a, b + 1), \text{ or } (2a, 2b). \quad (1.4)$$

This graph maps onto the graph (1.3) by $(a, b) \mapsto g^{ay+b}$, where y is the (secret and unknown) exponent of $h = g^y$.

Our argument has two main ingredients. The first is a spectral upper bound on the mixing time of this graph, which roughly speaking shows that the coefficients (a_k, b_k) become equidistributed after a small number of steps. This is very similar to the argument in [5] to guarantee collisions among the x_k . That alone, however, is not enough to show nondegeneracy: it is important to note that *undirected* 3-regular graphs can have this equidistribution feature, while simultaneously having (a_k, b_k) equaling (a_{k+2}, b_{k+2}) with probability $\geq 1/3$ (for example, going backwards on the edge just traveled). The second ingredient, an estimate on the number of short cycles, handles this. It is this part which depends on the condition on the multiplicative order of 2 modulo n , and hence which is not completely general.

We conclude this section with the proof of Theorem 1.2, which depends on estimates of the last two sections. Section 2, roughly speaking, deals with long random walks, while Section 3 with short random walks. The condition on the multiplicative order of 2 modulo n is needed to make sure their intervals of applicability overlap.

Proof of Theorem 1.2. Once a collision occurs, all future collisions are non-degenerate if and only if the first one was; this is because of the invertibility of the steps in (1.4). Thus, it suffices to assume that no collision has occurred until time T , which allows us to model the coefficients (a_k, b_k) , $k \leq T$, using a random walk on (1.4). Because the starting point $x_0 = g^{r_1} h^{r_2}$ is uniformly distributed and walk up to time T is random, the values of each (a_k, b_k) are themselves uniformly distributed. We show in Proposition 2.18 and Lemma 3.1 that for any $m > 0$ and a random point $(a, b) \in (\mathbb{Z}/n\mathbb{Z})^2$, a random walk of length m starting at (a, b) ends at (a, b) with probability at most $\frac{3}{2} \frac{1}{n^2}$. By the union bounds, the probability of a degeneracy $(a_k, b_k) = (a_\ell, b_\ell)$ occurring for some distinct $k, \ell \leq T$ is bounded above by

$$\sum_{\substack{k, \ell \leq T \\ k \neq \ell}} P[(a_k, b_k) = (a_\ell, b_\ell)] \leq \sum_{\substack{k, \ell \leq T \\ k \neq \ell}} \frac{3}{2} \frac{1}{n^2} < \frac{3}{2} \frac{T^2}{n^2}. \quad (1.5)$$

□

It is a pleasure to acknowledge Ravi Montenegro, Ze'ev Rudnick, Adi Shamir, and Prasad Tetali for their helpful discussions. We also thank Curt McMullen for helpful comments concerning the remark at the end of Section 2, and the anonymous referee for their suggestions for improving the paper, in particular Remark 3 above.

2 Mixing time estimates

Let A denote the adjacency operator of the graph (1.4): it is defined on complex-valued functions f on $(\mathbb{Z}/n\mathbb{Z})^2$ by the formula

$$Af(a, b) = f(2a, 2b) + f(a + 1, b) + f(a, b + 1). \quad (2.1)$$

Such functions themselves form a complex vector space of dimension n^2 , which is equipped with the usual inner product and norm

$$\langle f_1, f_2 \rangle = \sum_{a, b \in \mathbb{Z}/n\mathbb{Z}} f_1(a, b) \overline{f_2(a, b)}, \quad \|f\|^2 = \langle f, f \rangle. \quad (2.2)$$

Of special interest to us is the restriction of A to $\mathbb{1}^\perp$, the orthogonal complement of the constant function $\mathbb{1}$ (the functions on the graph whose average value is zero). The following result relates the operator norm properties of this restriction of A , to the mixing properties of the random walk on the graph:

2.3 Lemma. ([5, Lemma 2.1]) *Let Γ denote a directed graph on the vertex set V , having both d directed edges entering and exiting each vertex (including multiplicity). Suppose that there exists a constant $\mu < d$ such that $\|Af\| \leq \mu\|f\|$ for all $f \in \mathbb{1}^\perp$. Let S be an arbitrary subset of V . Then the number of paths of length $r \geq \frac{\log(2n)}{\log(d/\mu)}$ which start from any given vertex and end in S is between $\frac{1}{2}d^r \frac{|S|}{|V|}$ and $\frac{3}{2}d^r \frac{|S|}{|V|}$.*

Thus sufficiently long random walks hit a set S with probability between $\frac{1}{2} \frac{|S|}{|V|}$ and $\frac{3}{2} \frac{|S|}{|V|}$, independent of their starting point. Unfortunately this Lemma does not apply directly to our situation, because it can happen that $\|Af\| = \|f\|$ for some functions f . However, to show that random walks mix it suffices to work two steps at a time; fortunately, a nontrivial operator norm estimate

applies to A^2 instead, which corresponds to the adjacency operator for the graph on $(\mathbb{Z}/n\mathbb{Z})^2$ with edges

$$(a, b) \longrightarrow (4a, 4b), (2a+1, 2b), (2a, 2b+1), (2a+2, 2b), (2a, 2b+2), \\ (a+1, b+1), (a+1, b+1), (a+2, b), \text{ or } (a, b+2). \quad (2.4)$$

2.5 Proposition. *With A denoting the adjacency operator of the graph (1.4) and the standing assumption that n is an odd prime, there exists an absolute constant $c > 0$ such that*

$$\|A^2 f\| \leq \left(3 - \frac{c}{(\log n)^2}\right)^2 \|f\|, \quad f \in \mathbb{1}^\perp. \quad (2.6)$$

Proof. Any function f on the vertices may be expanded in terms of the additive characters $\chi_{k,\ell}(x, y) = e^{2\pi i(kx + \ell y)/n}$:

$$f = \sum_{k, \ell \in \mathbb{Z}/n\mathbb{Z}} c_{k,\ell} \chi_{k,\ell}. \quad (2.7)$$

The condition that $f \in \mathbb{1}^\perp$ is equivalent to $c_{0,0} = 0$. The action of A on the character $\chi_{k,\ell}$ is given by

$$A \chi_{k,\ell} = d_{k,\ell} \chi_{k,\ell} + \chi_{2k,2\ell}, \quad (2.8)$$

where

$$d_{k,\ell} = e^{2\pi i k/n} + e^{2\pi i \ell/n}. \quad (2.9)$$

Thus A is the sum of the diagonal operator $D : \chi_{k,\ell} \mapsto d_{k,\ell} \chi_{k,\ell}$ and the permutation operator $P : \chi_{k,\ell} \mapsto \chi_{2k,2\ell}$. The adjoint of A under the inner product (2.2) is $A^* = \overline{D} + P^{-1}$. Let us write

$$A^{*2} A^2 = (\overline{D}^2 + P^{-1} \overline{D} + \overline{D} P^{-1} + P^{-2})(D^2 + PD + DP + P^2) \\ = X_1 + X_2, \quad (2.10)$$

where $X_1 = \overline{D}^2 PD + \overline{D} P$, and X_2 is the remaining sum of 14 terms from the expansion of the first line. Because $|d_{k,\ell}| = 2|\cos(\frac{\pi(k-\ell)}{n})| \leq 2$ and in fact equals 2 when $k = \ell$, the operator norms of D and \overline{D} are $\|D\| = \|\overline{D}\| = 2$. Likewise $\|P\| = \|P^{-1}\| = 1$, because P preserves norms. It follows from the

sum of 14 terms defining X_2 that $\|X_2\| \leq 71$. Using this fact and Cauchy-Schwartz, we get the bound

$$\|A^2 f\|^2 = \langle f, A^{*2} A^2 f \rangle \leq 71 \|f\|^2 + |\langle f, (\overline{D}^2 PD + \overline{D}P)f \rangle|. \quad (2.11)$$

In order to prove (2.6) it now suffices to show the bound

$$\begin{aligned} |\langle f, (\overline{D}^2 PD + \overline{D}P)f \rangle| &\leq \left(10 - \frac{c}{(\log n)^2}\right) \|f\|^2 \\ &= \left(10 - \frac{c}{(\log n)^2}\right) n^2 \sum_{(k,\ell) \neq (0,0)} |c_{k,\ell}|^2, \end{aligned} \quad (2.12)$$

for some absolute constant $c > 0$. Here we have used (2.7) as well as the inner product relation

$$\langle \chi_{k,\ell}, \chi_{k',\ell'} \rangle = \begin{cases} n^2, & (k,\ell) = (k',\ell') \\ 0, & \text{otherwise.} \end{cases} \quad (2.13)$$

Since

$$(\overline{D}^2 PD + \overline{D}P)\chi_{k,\ell} = \mu_{k,\ell} \chi_{2k,2\ell}, \quad \mu_{k,\ell} = \overline{d_{2k,2\ell}}^2 d_{k,\ell} + \overline{d_{2k,2\ell}}, \quad (2.14)$$

we have that

$$|\langle f, (\overline{D}^2 PD + \overline{D}P)f \rangle| \leq n^2 \sum_{(k,\ell) \neq (0,0)} |c_{k,\ell}| |c_{2k,2\ell}| |\mu_{k,\ell}|. \quad (2.15)$$

We now group the indices $(k,\ell) \neq (0,0)$ into the $n+1$ lines through the origin in $(\mathbb{Z}/n\mathbb{Z})^2$. Using the bounds

$$\begin{aligned} |\mu_{k,\ell}| &\leq 8 + 2 \left| \cos \frac{2(k-\ell)\pi}{n} \right|, \quad \text{for the lines with } k \neq \ell, \\ |\mu_{k,k}| &\leq 6 + 4 \left| \cos \frac{k\pi}{n} \right|, \end{aligned} \quad (2.16)$$

and the fact that 2 is invertible modulo n , the desired bound (2.12) reduces to the estimates

$$\begin{aligned} \sum_{k=1}^{n-1} x_k x_{2k} &\leq \sum_{k=1}^{n-1} x_k^2 \quad \text{and} \\ \sum_{k=1}^{n-1} x_k x_{2k} \left| \cos \frac{\pi k}{n} \right| &\leq \left(1 - \frac{c'}{(\log n)^2}\right) \sum_{k=1}^{n-1} x_k^2, \quad \text{for some absolute } c' > 0, \end{aligned} \quad (2.17)$$

for any real numbers x_1, \dots, x_{n-1} . The first inequality follows from $x_k x_{2k} \leq \frac{1}{2}(x_k^2 + x_{2k}^2)$, and the second is proven in [5, Prop. 3.1]. \square

Combining these, we have shown

2.18 Proposition. *There exists an absolute positive constant c_0 such that the number of paths of length $r \geq c_0(\log n)^3$ on the graph (1.4) which begin at the vertex (a, b) and end at the vertex (a', b') is between $\frac{1}{2} \frac{3^r}{n^2}$ and $\frac{3}{2} \frac{3^r}{n^2}$. In particular, the probability of a random walk of length $r \geq c_0(\log n)^3$ ending at its starting point is at most $\frac{3}{2} \frac{1}{n^2}$.*

Remarks: This mixing time estimate for the random walk can also be proven using the method of canonical paths from [6].

Interestingly, the mixing time estimate for the Pollard ρ graph (1.3) in [5] does not need the step $x \rightarrow xh$: the steps $x \mapsto xg$ and $x \mapsto x^2$ suffice. This follows from the same method of proof, and is suggested by the heuristic that the $x \mapsto xh$ step is approximated by the other two. However, the Pollard ρ coefficient graph does not rapidly mix unless all three steps in (1.4) are present. In any case, all three steps are necessary for the execution of the Pollard ρ algorithm.

3 Trace estimates

3.1 Lemma. *If $k \geq 1$ is less than the multiplicative order of 2 modulo n , then there are precisely $3^k - 2^k$ closed cycles on the graph (1.4) of length k . In particular, if (a, b) is a random point in $(\mathbb{Z}/n\mathbb{Z})^2$, the probability is at most $\frac{1}{n^2}$ that a random walk of length k which starts at (a, b) also ends at (a, b) .*

The number of such cycles is also given by $\text{tr } A^k$, where A is the adjacency operator of the graph. The above estimate, however, does not seem to follow from the spectral techniques of the previous section.

Proof. Every path involves either doubling the coefficients (a, b) , or adding 1 to one of them. Thus all paths of length k starting from the vertex (x, y) have the form

$$T : (x, y) \mapsto 2^s(x, y) + (u, v), \quad (3.2)$$

where $s \leq k$ equals the number of doubling steps in the path, and $u, v \in \mathbb{Z}/n\mathbb{Z}$ are independent of x and y . (This characterization obviously holds for $k = 1$, and in general by induction.) A closed cycle is equivalent to a fixed point for T . Of the 3^k possible paths starting from (x, y) , exactly 2^k have $s = 0$. For those walks, $(u, v) \neq (0, 0)$ since all steps are of the form $(a, b) \mapsto (a + 1, b)$ or $(a, b + 1)$. Thus T has no fixed points in this situation.

However, if $s \neq 0$ then 2^s is not congruent to 1 modulo n because of the multiplicative order condition. In this situation, T has exactly one fixed point. The closed cycles come from these $3^k - 2^k$ cases. \square

This concludes the estimates necessary for the proof of Theorem 1.2. We conclude with the following lemma, which shows that primes for which 2 has multiplicative order smaller than its condition $c_0(\log n)^3$ are extremely rare.

3.3 Lemma. *Let $X > 0$, $c > 0$, and B the set of primes p in the interval $[X, 2X]$ such that the multiplicative order of 2 modulo p is bounded by $c(\log X)^3$. Then the size of B is bounded by*

$$|B| \leq \frac{c^2 \log(2)}{2} (\log X)^5. \quad (3.4)$$

Proof. The condition on $p \in B$ states that p divides $\prod_{k \leq c(\log X)^3} (2^k - 1)$. Primality therefore implies that

$$\prod_{p \in B} p \text{ divides } \prod_{k \leq c(\log X)^3} (2^k - 1), \quad (3.5)$$

and in particular satisfies

$$X^{|B|} \leq \prod_{p \in B} p \leq \prod_{k \leq c(\log X)^3} (2^k - 1) \leq 2^E, \quad (3.6)$$

with

$$E = \sum_{k \leq c(\log X)^3} k \leq \frac{1}{2} c^2 (\log X)^6, \quad (3.7)$$

implying (3.4). \square

In fact this proof shows something slightly stronger, that the bound (3.4) holds for the number of primes at least X whose multiplicative order is bounded by $c(\log X)^3$.

References

- [1] Stephen C. Pohlig and Martin E. Hellman, *An improved algorithm for computing logarithms over $\text{GF}(p)$ and its cryptographic significance*, IEEE Trans. Information Theory **IT-24** (1978), no. 1, 106–110.

- [2] Jeremy Horwitz and Ramarathnam Venkatesan, *Random Cayley digraphs and the discrete logarithm*, Algorithmic number theory (Sydney, 2002), 2002, pp. 416–430.
- [3] J.-H. Kim, R. Montenegro, and P. Tetali, *Near Optimal Bounds for Collision in Pollard Rho for Discrete Log*, Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2007), 2007, pp. 215–223.
- [4] J.-H. Kim, R. Montenegro, Y. Peres, and P. Tetali, *A Birthday Paradox for Markov chains, with an optimal bound for collision in Pollard Rho for Discrete Logarithm*, Algorithmic number theory, Lecture Notes in Comput. Sci., vol. 5011, Springer, Berlin, 2008, pp. 402–415.
- [5] Stephen D. Miller and Ramarathnam Venkatesan, *Spectral analysis of Pollard rho collisions*, Algorithmic number theory, Lecture Notes in Comput. Sci., vol. 4076, Springer, Berlin, 2006, pp. 573–581.
- [6] Ravi Montenegro and Prasad Tetali, *Mathematical Aspects of Mixing Times in Markov Chains*, Foundations and Trends in Theoretical Computer Science, 2006.
- [7] Victor Shoup, *Lower bounds for discrete logarithms and related problems*, Advances in cryptology—EUROCRYPT '97 (Konstanz), 1997, pp. 256–266. Updated version at <http://www.shoup.net/papers/dlbounds1.pdf>.

Addresses:

Stephen D. Miller
 Department of Mathematics
 110 Frelinghuysen Road
 Rutgers, The State University of New Jersey
 Piscataway, NJ 08854
miller@math.rutgers.edu

Ramarathnam Venkatesan
 Microsoft Research Cryptography and Anti-Piracy Group
 1 Microsoft Way
 Redmond, WA 98052
 and
 Cryptography, Security and Applied Mathematics Group
 Microsoft Research India
 Scientia - 196/36 2nd Main
 Sadashivnagar, Bangalore 560 080, India
venkie@microsoft.com